arXiv:1412.0003v1 [cs.CV] 26 Nov 2014

3D-Assisted Image Feature Synthesis for Novel Views of an Object

Hao Su* Fan Wang* Li Yi Leonidas Guibas
Stanford University

December 2, 2014

Abstract

Comparing two images in a view-invariant way has been a challenging problem in computer vision for a long time, as visual features are not stable under large view point changes. In this paper, given a single input image of an object, we synthesize new features for other views of the same object. To accomplish this, we introduce an aligned set of 3D models in the same class as the input object image. Each 3D model is represented by a set of views, and we study the correlation of image patches between different views, seeking what we call surrogates — patches in one view whose feature content predicts well the features of a patch in another view. In particular, for each patch in the novel desired view, we seek surrogates from the observed view of the given image. For a given surrogate, we predict that surrogate using linear combination of the corresponding patches of the 3D model views, learn the coefficients, and then transfer these coefficients on a per patch basis to synthesize the features of the patch in the novel view. In this way we can create feature sets for all views of the latent object, providing us a multi-view representation of the object. View-invariant object comparisons are achieved simply by computing the L^2 distances between the features of corresponding views. We provide theoretical and empirical analysis of the feature synthesis process, and evaluate the proposed view-agnostic distance (VAD) in fine-grained image retrieval (100 object classes) and classification tasks. Experimental results show that our synthesized features do enable view-independent comparison between images and perform significantly better than traditional image features in this respect.

1 Introduction and Related Work

Object recognition plays a key role in many applications and has been a central topic in the vision community. The fundamental question — easy to state, albeit harder to formalize — is: when do two or more images are said to be "same" or "different"? Appearance differences of images can be factored as intrinsic or extrinsic. Intrinsic factors refer to properties of the imaged objects themselves, such as differences in the

*indicates equal contributions.

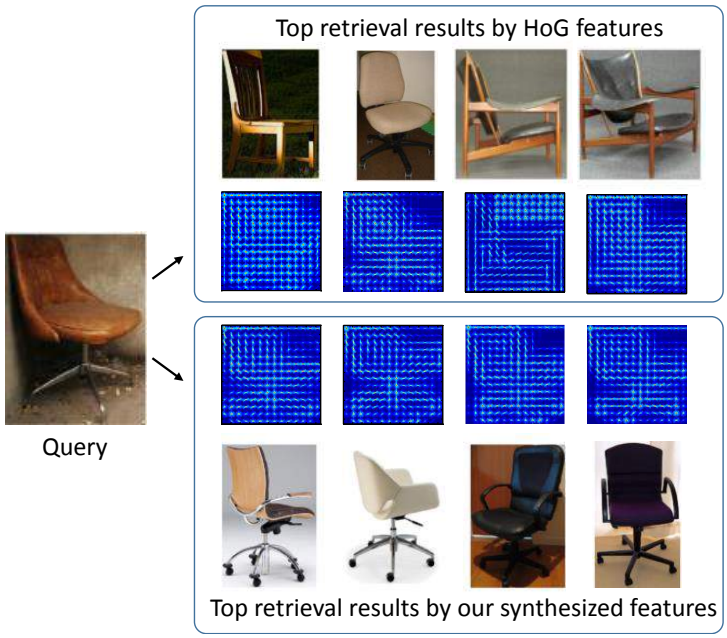

Figure 1: **View-invariant image retrieval.** Given each image, we synthesize its image features on a predefined list of viewpoints, and compare two images by their synthesized features from all viewpoints. Therefore, we can find images with similar objects regardless of viewpoint change. The first row shows the retrieval results by the HoG features on the original images only, and the second row visualizes the HoG features used for retrieval. The third row visualizes the synthesized HoG features, and the last row shows retrieval results using our synthesized features. Synthesized features are shown only on the view corresponding to the query image due to paper length constraint.

topology, geometry and material of the underlying or latent imaged shapes. Extrinsic factors on the other hand include illumination, viewpoint, and more generally visibility effects (occlusion, etc.). For object recognition and image understanding, it is a fundamental problem to find the "best" representation which only focuses on intrinsic object properties and is invariant to extrinsic factors (Fig. 1).

In the vision literature, substantial efforts have been made to achieve such invariant representations. Many kinds of view-invariant image features have been proposed for image comparison and recognition. The basic idea is to embed images into a *common space*, so that the embedded point remains relatively stable as the viewpoint changes. The common space is usually built out of low-level image features with careful engineering [13, 3, 15]. Recently, there have also been successful efforts that learn such view-invariant features from large image datasets [10, 12]. However, these methods can only achieve invariance for small viewpoint changes, typically not more than 15°. Another option is to choose the common space to be a high-level conceptual representation, based on object class or attributes [5, 11]. However, much detailed geometric and physical information about the object gets lost in such embedding pro-

cesses. There has also been work on reconstructing 3D geometry from an image [19]; however, these algorithms still lack the ability to recover detailed information and do not scale well.

In this paper, we choose the common space to be the *space of 3D shapes*, similar to [19]. Objects in images are indeed 2D projections of *latent* 3D shapes; therefore, if we can obtain the latent 3D shape for each image, then images can be compared in a view-agnostic manner. Compared with low-level features, this approach achieves view-invariance for arbitrary viewpoint changes. Compared with high-level representations, 3D shape space is closer to the physical form of objects and therefore preserves more basic and detailed information.

Reconstructing 3D geometry, however, is a really challenging task and unnecessary in many cases, for example, image comparisons are typically based on image feature sets. In this paper *we focus on reconstructing the features of different views of the latent shape of an imaged object* (Fig. 1). So effectively we choose a multi-view representation of a shape, also known as 2-1/2D representation [18, 4], where each 3D shape is described by a set of images rendered from a predefined list of viewpoints. The descriptor of the latent shape is just a concatenation of image features from each of the views (see §3.1). In this way, the problem of reconstructing the features of the latent shape can be formulated as: given an object image (one view), reconstruct its features from other novel views (in the desired view set).

This problem is very challenging, because it is naturally ill-posed. The input is only one single real image based on one view — thus information seems to be missing for reconstruction from novel views, even if all we seek is features for the new views. Therefore, we introduce a 3D shape collection from the same object class as a non-parametric prior. The intuition of our proposed method is that, given the novel view, we find related parts in the observed view that can best help us estimate the novel view. For this task, we have two guides: the image from the observed view, and the entire shape collection. Thus we explore two kinds of relationships to accomplish feature synthesis.

The first type of relationship reflects the *intra-shape structure* which builds the relationships between the novel view and the observed view. More specifically, such relationships characterize the correlation of features at different locations of different views. Such correlations naturally exist because images from different views observe the same underlying 3D shape, whose parts may be further correlated by 3D symmetries, repetitions, and other factors. We use a probabilistic framework to quantitatively measure such correlations, aiming to estimate the "surrogate suitability" of one image patch in one view to predict another patch in another view. Such relationships can be discovered efficiently and accurately from the shape collection.

The second type of relationship reflects the *inter-shape structure* which builds the relationships between the image object and the shape collection. Although the entire shape space for our multi-view shape representation is highly nonlinear, local neighborhoods can be well approximated by a linear low-dimensional subspace [20]. This allows us to synthesize novel points in the shape space through linear interpolation, so as to approximate the latent image object. The key point for capturing this relationship is to estimate appropriate coefficients for the interpolation, and we use an approach derived from locally linear embedding (LLE) methods [16].

To summarize, for each patch in the novel view, the intra-shape relationships allows us to find which patches in the observed view are its best surrogates, and the inter-shape relationships teach us how the feature of the new patch should be synthesized from those of its surrogates. In this way we can populate with features for all views of the latent object in our image, effectively creating its representation in our shape space.

Our major contributions in this paper are:

- We propose a method for synthesizing object image features from unobserved views by exploiting inter-shape and intra-shape relationships;
- Given the synthesized image features for novel views, we are able to compare two images of the same or different objects by comparing their synthesized multi-view shape features. The resulting distance is view-invariant and achieves much better performance on fine-grained image retrieval and classification tasks when compared with previous methods.

2 Problem Formulation and Method Overview

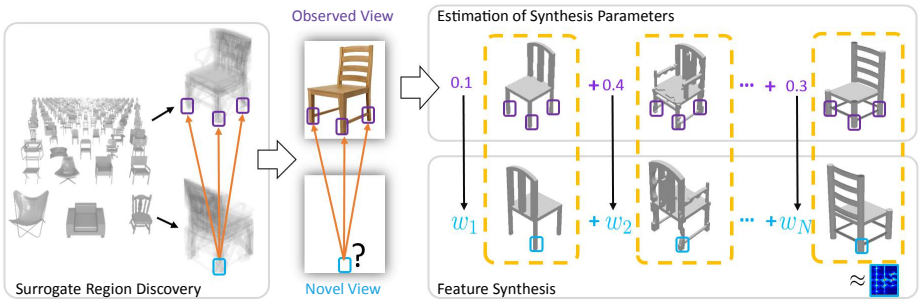

Figure 2: **Method overview.** Given an object image, we synthesize image features for novel views of the latent underlying object. The synthesis is done patch-by-patch. To predict the feature in the blue patch, we first look for regions in the observed view that are the most correlated to it, denoted as the surrogate regions (purple patches). The surrogate regions are found by analyzing the shape collection (**Surrogate Region Discovery**, §3.2). Then, at the observed view, we learn how to reconstruct the surrogate region by a linear combination of the same region in the same view from all shapes in the shape collection (**Estimation of Synthesis Parameter**, §3.3). Finally, we transfer the linear combination coefficients back to the novel view to reconstruct the features in the blue patch by a linear combination of the features at the same patch on the novel view from all shapes (**Feature Synthesis**, §3.4).

Problem Input Our input contains two parts:
1) an image of an object O with bounding box and known class label. With recent advances in image detection and classification [17], obtaining object label and bounding box has become much easier than before. All following steps are performed on a cropped image which only contains the object.

2) a collection of 3D shapes (CAD models) from the same class. All 3D shapes are orientation aligned in the world coordinate system during a preprocessing step. Each shape is stored as a group of rendered images from the predefined list of viewpoints. Each rendered image is also cropped around the object. The view for object O in the input image is estimated and approximated by one of the predefined viewpoints (§4.1). To preserve detail information, the input image and the rendered images are resized to a fixed size and partitioned into overlapping square patches. Patch-based features such as HoG are extracted for each patch.

Problem Output The output is the multi-view shape representation of the *latent shape* of the input object, consisting of one image descriptor for each of the predefined views.

Without loss of generality, the key subproblem can be formulated as: given the object image in the input viewpoint v_0, estimate its features from another novel viewpoint v_1. The full multi-view representation can then be obtained by repeating this process for each predefined viewpoint.

Method Overview The proposed framework is shown in Fig. 2. For a specific patch in the novel view (the query patch), we seek to find the patches on the observed view which can best predict it (*Surrogate Region Discovery* in Fig. 2), and then learn how the features in those "surrogate" patches at the observed view can be best synthesized by the 3D model views (*Estimation of Synthesis Parameters* in Fig. 2). We finally apply the same synthesis method to the desired query patch (*Feature Synthesis* in Fig. 2).

3 Novel View Image Feature Synthesis

3.1 Notation

Our notation follows standard mathematical conventions. The set of preselected viewpoints is indexed by $\mathcal{V} = \{1, \ldots, V\}$. Each rendered image or the input real image is covered by G overlapping patches, indexed by $\mathcal{G} = \{1, \ldots, G\}$. A patch-based feature set $\mathbf{f} = [\mathbf{x}_1^T; \ldots; \mathbf{x}_G^T] \in \mathbb{R}^{G \times D}$ is extracted for the image, where each $\mathbf{x}_g \in \mathbb{R}^D$ is a feature vector for patch g. So the multi-view shape descriptor is represented by a tensor $\mathbf{S} = [\mathbf{f}_1; \ldots; \mathbf{f}_V] \in \mathbb{R}^{V \times G \times D}$, in which each $\mathbf{f}_v$ is a feature of a rendered image at view v. Finally, the 3D shape collection is denoted by $\mathcal{S} = \{\mathbf{S}_1, \ldots, \mathbf{S}_N\}$, where $\mathbf{S}_n$ denotes the multi-view descriptor of a shape n. For convenience, we further let $\mathbf{S}_{n,v,g} \in \mathbb{R}^D$ denote the features of the g-th patch in the v-th view of the n-th shape.

3.2 Surrogate Region Discovery

Our ultimate goal is to transfer information across views, since we want to apply the synthesis parameter learned from one patch on one view to some other patches on other views. Therefore, we first exploit the cross-view patch appearance correlation as an important building block.

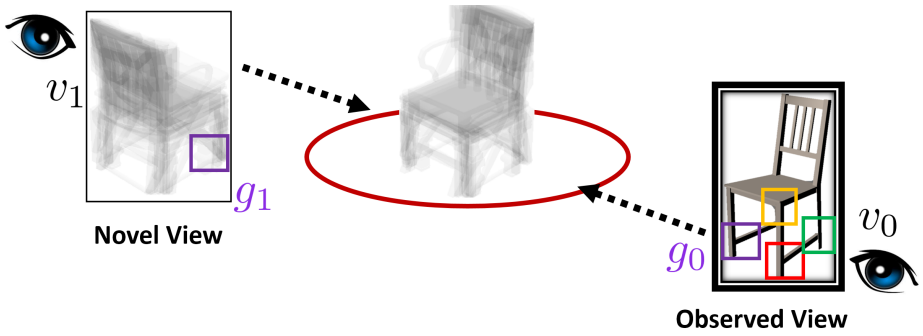

Figure 3: **Patch surrogate relationship** (§3.2)

Fig. 3 shows some intuitive examples about patch relationships. It is obvious that observing patch g_0 at view v_0 (purple box) helps us to determine the appearance of patch g_1 at view v_1 because they correspond to the same leg of a chair. Besides that, other factors such as symmetry and part membership in 3D shapes can also induce strong correlations among patch appearances. For example, the red patch in v_0 strongly correlates with g_1 because of chair symmetry; the green patch in v_0 is also correlated with g_1 because it belongs to the same part type as g_1 (chair leg). On the other hand, the appearance of chair back at v_0 will not be very helpful in determining g_1.

Therefore, there exists a group of patches at the observed view which are correlated with the query patch at the novel view, which we call *surrogate patches*; the region they form is called a surrogate region $\mathcal{R}$.

This relationship between patches across views can possibly be inferred by analyzing the shape geometry, but this is non-trivial and would require reliable object part segmentation, symmetry detection,etc. Therefore, we introduce an learning-based approach instead.

To precisely quantify such correlations between patches, we first introduce the concept of **perfect patch surrogate**:

Definition 1. *Patch g_0 at view v_0 is a* **perfect patch surrogate** *for patch g_1 at view v_1 if $\mathbf{S}_{i,v_0,g_0} = \mathbf{S}_{j,v_0,g_0}$ implies $\mathbf{S}_{i,v_1,g_1} = \mathbf{S}_{j,v_1,g_1}$ for* **any** *shape pair $\mathbf{S}_i$ and $\mathbf{S}_j$.*

Intuitively, this means the similarity of patch g_0 at view v_0 implies the similarity of patch g_1 at view v_1 between a pair of 3D shapes. Usually patches cannot be perfect surrogates for each other, so we seek for a probabilistic version of Definition 1:

Definition 2. *For a given patch g_1 at v_1, the* **surrogate suitability** *of patch g_0 at view v_0 is defined as*

$$\gamma(g_0; g_1) = \log P(\mathbf{S}_{i,v_1,g_1} = \mathbf{S}_{j,v_1,g_1} | \mathbf{S}_{i,v_0,g_0} = \mathbf{S}_{j,v_0,g_0}),$$

$\gamma(g_0; g_1)$ is a measure of how suitable patch g_0 is as a surrogate for patch g_1. Intuitively, larger $\gamma(g_0; g_1)$ indicates a stronger correlation (Fig. 4). Therefore, the surrogate region $\mathcal{R}(g_1)$ can either consist of the top k_p patches with highest $\gamma(g_0; g_1)$, or $\mathcal{R}(g_1) = \{g_0 : \gamma(g_0; g_1) > \tau\}$, where k_p or τ is determined empirically. We discuss the estimation of $\gamma(g_0, g_1)$ in the next section.

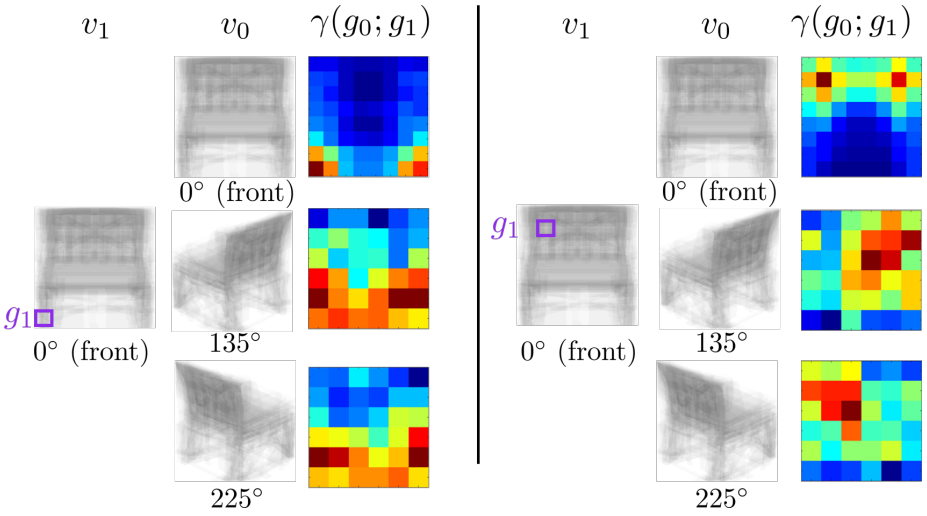

Figure 4: **Visualization of patch surrogate suitability.** Two examples of the surrogate suitability from g_1 in v_1 to patches in view v_0. Red means large γ. For example, in the left figure, g_1 corresponds to the tip of right-front leg at v_1 (front view). At the front view itself, the left-front and right-front leg tips have higher surrogate suitability for g_1 because of symmetry; at the 225° view, the left-back, right-back and right-front leg tips have higher surrogate suitability because of symmetry and part membership.

3.2.1 Estimation of Patch Surrogate Suitability

With the large-scale shape collection at hand, we adopt a learning based approach to estimate the (probabilistic) patch surrogate suitability in a data-driven manner.

Estimating $\gamma(g_0; g_1)$ is a non-parametric density estimation problem. As image features are high-dimensional continuous variables, theoretical results indicate that the sample complexity for reliable estimation is very high and infeasible in practice. To overcome the difficulty, we quantize features into a vocabulary $\mathcal{D}$ containing D visual words. For notation convenience, we denote the codeword of $\mathbf{S}_{i,v_0,g_0}$ by $A^i_{g_0}$ and $\mathbf{S}_{i,v_1,g_1}$ by $A^i_{g_1}$, then

$$\gamma(g_0; g_1) = \log P(A^i_{g_1} = A^j_{g_1} | A^i_{g_0} = A^j_{g_0}) \tag{1}$$

where P is the probability measure.

Estimating (1) by an empirical conditional distribution still requires a large amount of samples. However, we show that (1) can be cast as a Rényi entropy estimation problem. We can prove that the *optimal* sample complexity needed for estimating (1) is $\Theta(D)$ (Theorem 1 in Appendix). Roughly speaking, with $N = \Theta(D)$ shapes, we can accurately estimate (1) with high probability. The proof also suggests an algorithm

to estimate Eq (1) as below:

$$\hat{\gamma}(g_0; g_1) = \log \sum_{\substack{(A_{g_0}, A_{g_1}) \\ \in \mathcal{D} \times \mathcal{D}}} \hat{P}^2(A_{g_0}, A_{g_1}) - \log \sum_{A_{g_1} \in \mathcal{D}} \hat{P}^2(A_{g_1}) \,.$$

Here, probabilities $P^2(x)$ should be estimated by $\hat{P}^2(x) = \frac{N_x(N_x-1)}{N^2}$, where N_x is the total number of times value x appears in samples and $N = \sum_x N_x$.

3.3 Estimation of Synthesis Parameters

As we have mentioned, the global shape space for our multi-view representation is non-linear and high-dimensional. Our assumption, however, is that shapes in a local neighborhood can be well approximated by a locally linear and low-dimensional subspace. Since the multi-view representation is actually a concatenation of features from all patches of all views, this local linearity does not only hold for the whole shape, but it also holds for each view of the shape, for each patch of the view, or even for a subset of patches of the view. In other words, features for the patches from the same location(s) on the same view of all shapes also lie in a locally linear subspace.

For any patch g in view v, its feature is denoted as $\mathbf{x}_{v,g} \in \mathbb{R}^D$. We use $\mathbf{S}_{:,v,g} \in \mathbb{R}^{D \times N}$ to denote the feature matrix collecting patch g of view v of all 3D models, then local linearity tells us that

$$\mathbf{x}_{v,g} \approx \mathbf{S}_{:,v,g} \mathbf{w}_{v,g} \,, \tag{2}$$

where $\mathbf{w}_{v,g} \in \mathbb{R}^N$ is the reconstruction coefficient.

Given a surrogate region $\mathcal{R}$ on the observed view, its features should be a linear combination of the same region across different 3D shapes. So $\mathbf{w}_{v_0,\mathcal{R}}$ can be estimated by solving an Locally Linear Embedding (LLE) problem:

$$\begin{aligned} \underset{\mathbf{w}_{v_0,\mathcal{R}}}{\text{minimize}} \quad & \sum_{g_0 \in \mathcal{R}} \left\| \mathbf{x}_{v_0,g_0} - \mathbf{S}_{\mathcal{N},v_0,g_0} \mathbf{w}_{v_0,\mathcal{R}} \right\|^2 , \\ \text{subject to} \quad & \mathbf{w}_{v_0,\mathcal{R}} \geq \mathbf{0}; \quad \mathbf{w}_{v_0,\mathcal{R}}^T \mathbf{1} = 1 \,, \end{aligned} \tag{3}$$

where $\mathcal{N}$ denotes the k-nearest shapes obtained from the whole shape collection by comparing the rendered images on v_0 with the input image, thus $\mathbf{S}_{\mathcal{N},v_0,g_0} \in \mathbb{R}^{D \times k}$ and $\mathbf{w}_{v_0,\mathcal{R}} \in \mathbb{R}^k$.

Note that our reconstruction coefficient $\mathbf{w}_{v_0,\mathcal{R}}$ is specific to the choice of view v_0 and patch(es) $\mathcal{R}$, unlike previous locally linear reconstruction methods assuming uniform $\mathbf{w}$ for the whole image descriptor [21]. Experiments show that spatial-varying coefficients allow us to recover features more precisely and with better spatial locality as compared with image-wide uniform coefficients (Fig. 8c and Fig. 8d).

3.4 Feature Synthesis

Now that we have the synthesis coefficients estimated for $\mathcal{R}$ on view v_0, we have to decide how to transfer it back to v_1, so that we can synthesize $\mathbf{x}_{v_1,g_1}$ by apply the coefficients on features of g_1 on v_1 from all shapes.

We make the following assumption to connect the weight across views: for a shape $\mathbf{S}$, if a patch g_0 can surrogate g_1 very well (with high $\gamma(g_0; g_1)$), and $\mathbf{x}_{v_0,g_0} \approx \mathbf{S}_{:,v_0,g_0}\mathbf{w}_{v_0,g_0}$ and $\mathbf{x}_{v_1,g_1} \approx \mathbf{S}_{:,v_1,g_1}\mathbf{w}_{v_1,g_1}$, then their weights are the same, i.e. $\mathbf{w}_{v_0,g_0} \equiv \mathbf{w}_{v_1,g_1}$.

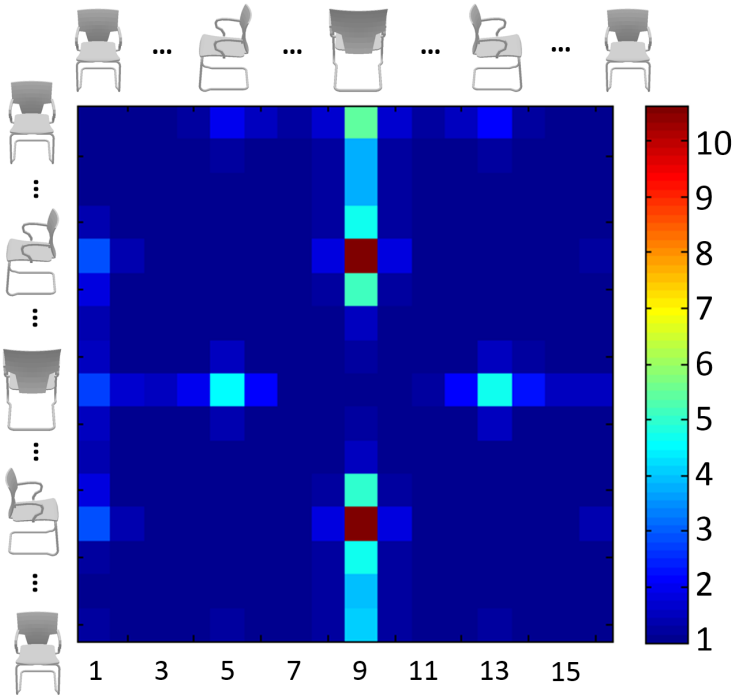

Figure 5: **Evaluation of weight transferability.** The experiments are performed on rendered images of the shape collection (see §4.3 for details of experiment setup). The (i, j)-th element is obtained as below: we estimate the coefficients on view v_i for all other shapes to reconstruct S_{0,v_i}, and apply them to v_j to obtain the reconstructed feature $\hat{S}_{0,v_j}$; then we check the ranking of the distance between ground truth S_{0,v_j} and reconstructed $\hat{S}_{0,v_j}$ in the list of distance values between S_{0,v_j} and all other shapes on v_j. (i, j)-th element shows the ranking averaged through the whole collection. If the ranking is 1, it means the reconstructed feature is the closest to the ground truth feature, indicating that the reconstruction is good enough to describe S_0 on v_j.

Empirical verification of this assumption is shown in Fig. 5. The (i, j)-th element in the matrix shows the transferability from view v_i to v_j. It measures how close the synthesized feature on view v_j is to its ground truth version when using coefficients estimated on v_i. Each entry could range from 1 to the size of model collection (5,057 in this experiment). The closer the value is to 1, the better the transferability is between v_i and v_j. The average value of the whole matrix is 1.39, quite close to 1, meaning that the weights transferred across views can reconstruct the features very well. Note that there are some entries indicating bad transferability between specific views. For example, view 5 and 9, which are the side view and back view respectively, cannot be transferred to each other very well because they share less common information.

Therefore, $\mathbf{w}_{v_1,g_1}$ can be replaced by $\mathbf{w}_{v_0,\mathcal{R}}$ if $\mathcal{R}$ is the appropriate surrogate region on v_0 for g_1. We can reconstruct the feature by $\mathbf{x}_{v_1,g_1} = \mathbf{S}_{\mathcal{N},v_1,g_1}\mathbf{w}_{v_0,\mathcal{R}}$. Fig. 6 shows two examples of our synthesized image features.

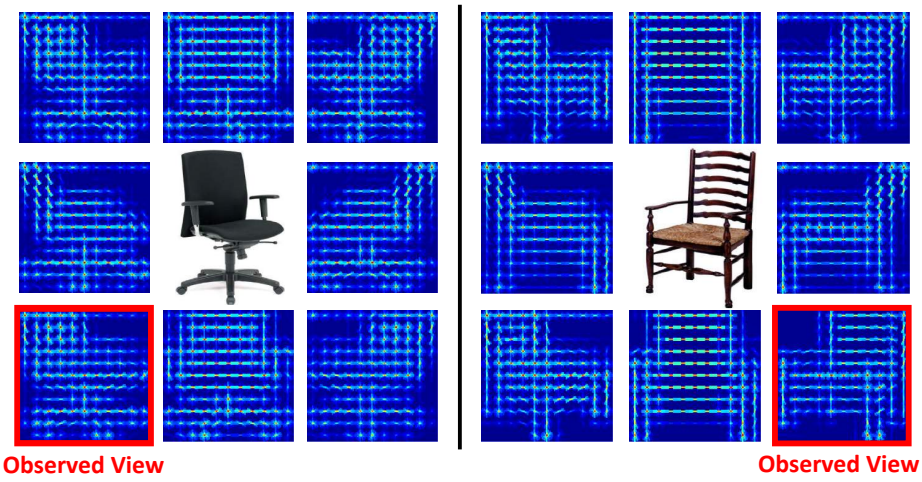

Figure 6: **Visualization of synthesized HoG features on 8 canonical views** On observed view, the original HoG feature is shown.

4 Experiments

4.1 Data Preparation

Large-scale Shape Dataset We introduce a large-scale shape collection containing human-built 3D meshes from 100 man-made object categories. There are ∼100,000 3D models in total. The number of models per class varies from 20 (purse) to over 13000 (table). The models of our dataset are from the Trimble 3D Warehouse and the annotations are crowd-sourced from Amazon Mechanical Turk (AMT). Each object category in our shape dataset is further mapped to a synset in ImageNet [6]. Please refer to the supplementary material for more details on the dataset.

Shape Collection Preprocessing To align the input 3D models, we employ the method described in [8], which jointly optimizes the orientation of all input 3D models to minimize the sum of distances between corresponding points computed using pair-wise alignment. To render 3D models, we sample a pre-specified number of different viewpoints over the viewing sphere centered at the shape.

The default setting for underlying features is as below unless specified otherwise: each rendered image is resized to 112×112 and partitioned into 32×32 patches which overlap with each other by 16 pixels, forming 6×6 patches in total; HoG features are extracted for each patch.

Image Preprocessing As noted in §2, we assume object bounding boxes are provided. Camera pose is estimated for the cropped object. We train a random forest using the rendered images of our aligned 3D models. Image features are extracted under default setting as before.

4.2 View-Agnostic Distance

Given the synthesized image features on a list of novel views, the distance between two images is obtained as the L2 distance between the two aligned and concatenated vectors of their synthesized multi-view features. Since the viewpoint information has been factored out, the resulting distance is view-invariant. This distance is denoted as view-agnostic distance (VAD) in the following experiments.

4.3 Method Analysis

In this section, we analyze the performance of our method under different settings and provide a thorough understanding. The quantitative analysis is done in a fine-grained image retrieval application. We take the class "Chair" as an example. We collect all images with bounding boxes in "Chair" and its fine-grained sub-categories from ImageNet, and verify their labels by AMT. There are in total 5,813 images in 15 fine-grained categories, denoted as a "cluttered" set. In contrast, a subset of 1,309 images with simple background is selected to form a "clean" set to help us better understand the performance. The "Chair" shape collection contains 5,057 models, rendered in 16 views.

In our experiment, each image is taken as query and all other images are sorted according to their distance to the query image. Images belonging to the same fine-grained category are regarded as correct. If the query image belongs to multiple fine-grain categories, returned images belonging to any of its categories are regarded as correct. Precision-recall curves are generated to evaluate the retrieval performance. Our proposed VAD is compared with L2 distance of the baseline HoG descriptors.

Fully Controlled Setting We first run an experiment using the data set from [2], which consists of 1393 3D chair models rendered in 62 viewpoints in an almost photo-realistic manner. We take all $1393 \times 62 = 86,366$ images and perform image retrieval with each image as query. Other view images from the same model are considered correct. This experiment is under fully controlled setting: the background is absolutely clean; pose estimation is 100% accurate, and there is even no pose discretization error, i.e. the estimated pose and the true pose of the imaged object are exactly the same. This setting is the most favorable for our proposed method; therefore, we achieve almost perfect performance as shown in Fig. 7a. Since the chair models in [2] and our models are from the same source, the chair models used here have been screened to avoid overlap with the 1,393 chairs, So it is guaranteed that there is no exact model for any of the query images.

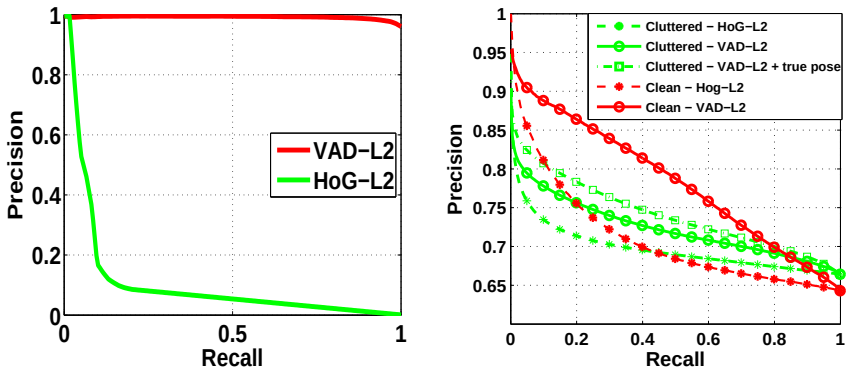

(a) Fully controlled setting.

(b) Clean vs. cluttered

Figure 7: **Performance comparison.**

Clean background v.s. Cluttered background Fig. 7b shows the precision-recall curves of HoG-L2 distance and the proposed VAD. In both the "clean" (red lines) and "clutter" (green lines) cases, VAD has greatly boosted the performance of baseline HoG-L2. However, in "clean" case, pose estimation and nearest 3D model matching are more accurate, thus the performance boost is more significant compared with "cluttered" case. With better pose estimation algorithm, our VAD still has space to improve. As shown in the green dash-dot line in Fig. 7b, we can boost the performance even more in cluttered case if the ground truth viewpoint of the input image is given.

Parameter Sensitivity Analysis Fig. 8a shows the performance changing with different number of 3D models. The performance for each experiment is summarized by the area under the precision-recall curve (AUC). The intuitive explanation is that, a larger shape collection is preferred since it can provide better coverage of the shape space and further help better reconstruct the descriptor on novel views. However, we also observe that the performance with 200 3D models is only 2% lower than the performance with the full collection of 5,057 3D models. The reason is that our model has the ability to "interpolate" in the shape space, which compensates for the absence of large shape collection at query time. We use the whole shape collection below.

Fig. 8b shows the performance changing with the parameter k for obtaining the local neighborhood in Eq (3). Specifically, for $k = 1$, it is equivalent to using the most similar shape to represent the query object. It is beneficial to use an appropriate range of neighborhood to find reconstruction coefficients of the query latent shape, which is shown in Fig. 8b. $k = 200$ is used for the rest experiments.

When finding the surrogate patch region on the observed view, the scope of locality can be defined by the number of patches k_p, or the threshold of the correlation score τ(§3.2). We show the performance changing with k_p or τ in Fig. 8c and Fig. 8d, respectively. $k_p = 36$ in Fig. 8c (the last data point) or $\tau = 0$ in Fig. 8d (the first data point) corresponds to the case when the whole image is selected as surrogate

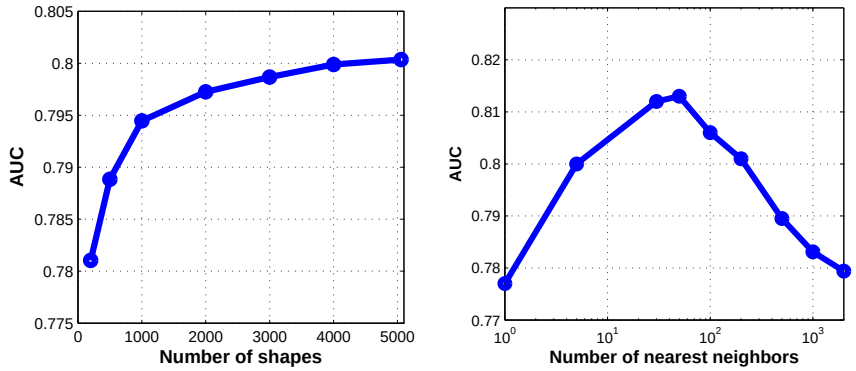

(a) Size of shape collection.

(b) Neighborhood for LLE

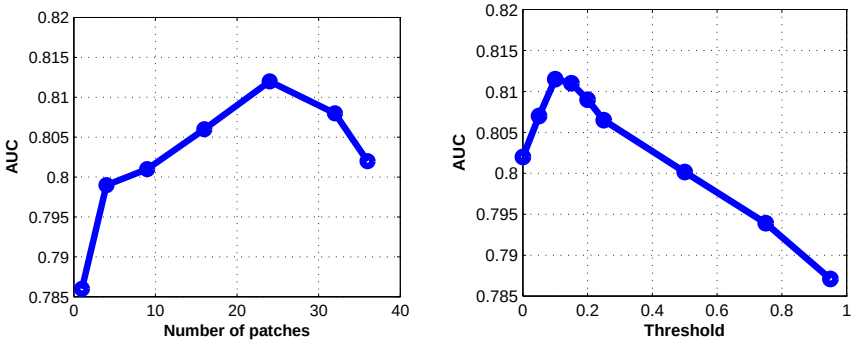

(c) Locality by # patches

(d) Locality by threshold

Figure 8: **Parameter sensitivity.**

region, i.e. LLE coefficients are uniform across the whole image. We can see that, an optimal region exists as a surrogate for a query patch. Including more surrogate patches can increase the samples for linear coefficients estimation, hence achieves better robustness; but including too many patches will eventually bring in unrelated patches, which are harmful.

We also investigate performance changing with k_p in different partitioning of image patches (Table 1). Similar conclusion can be made that the surrogate region cannot be too large nor too small. $k_p = 9$ for 6×6 is used as default.

Applicability for CNN Our approach is not restricted to any specific kind of descriptors. Features extracted by convolutional neural networks (CaffeNet [9]) from different layers are tested here to replace the HoG features. The performance is shown in Table 2. It can be seen that for different choices of underlying features, our method can always boost the performance.

	single	25%	50%	75%	100%
6x6	0.786	0.801	0.806	**0.81**	0.808
8x8	0.774	0.79	0.795	**0.797**	0.788
10x10	0.76	0.784	**0.787**	0.785	0.774

Table 1: **AUC for different patch partition and different number of patches for surrogate region.**

	Vanilla L2	VAD (ours)
CNN (CaffeNet, pool 1)	0.662	**0.677**
CNN (CaffeNet, pool 2)	0.697	**0.745**
CNN (CaffeNet, pool 5)	0.690	**0.746**
CNN (CaffeNet, fc 6)	0.748	**0.788**
CNN (CaffeNet, fc 7)	0.744	**0.785**

Table 2: **Performance by other kinds of image features.**

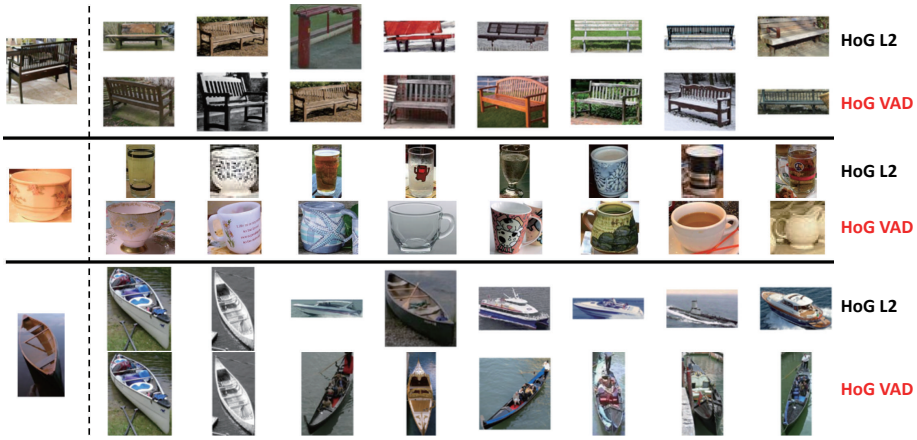

Figure 9: **Fine-grained retrieval results for "Bench", "Cup" and "Watercraft".**

4.4 Applications

Fine-grained Image Retrieval on 100 Classes We collect images of 100 classes with bounding boxes from ImageNet and verify their fine-grained labels using AMT. We have also preprocessed shape collections of the corresponding classes. Fine-grained image retrieval performance is evaluated as described in §4.3. On average, the baseline L2 distance of HoG descriptor can achieve AUC of 0.635, and our approach can achieve an AUC of 0.694. Fig. 9 shows some examples of retrieval results for comparison. More results can be found in appendix.

Part-based Image Retrieval Our approach can enable a new application of part-based image retrieval. The user can specify a region on the query image, and our approach can synthesize the features of related patches on novel views. The distance

between images will only be evaluated on these patches instead of the whole images. Fig. 10 shows examples of part-based image retrieval. The rectangles on query images are the input specified by users. Although the algorithm can only see the provided patch on the view of query image, it returns images with similar appearance in the corresponding regions from other viewpoints. This part-based search can be useful in product search by image, allowing users to express preferences for product parts.

Fine-Grained Object Categorization Besides image retrieval, we also evaluate our proposed distance on fine-grained object categorization. We use the FGVC-aircraft dataset [14], which contains 10,000 images with 100 different aircraft model variants. We use the non-linear SVM on a χ^2 kernel and replicate the SPM feature setting in [14], i.e, 600 k-means bag-of-visual words dictionary, multi-scale dense SIFT features, and 1×1, 2×2 spatial pyramid features. Using our approach, we obtain the view-invariant version of SPM feature. We also use bounding boxes predicted by R-CNN [7] and random forest for pose estimation (§4.1) on test data. Table 3 shows that our method significantly outperforms the baseline. Note that the baseline method in [14] does not use object bounding boxes in testing. To be fair, we also provide the baseline performance with bounding boxes provided.

Figure 10: **Part-based retrieval results**

5 Conclusion and Future Work

In this paper, we have proposed a framework for synthesizing features of a new view of an object in an image, given a collection of 3D models from the same object class. By collecting together the features from several canonical views of the object, we arrive at a view-independent model of the object. With this representation, we can achieve view-invariant image comparison, factoring out the influence of viewpoint and only focusing

	[14] (SPM)	[14] with b.box	Ours
Accuracy	0.487	0.561	0.603

Table 3: **Accuracy comparison on FGVC-aircraft.** Note that our results is based on [14] with bounding boxes.

on the intrinsic object properties. The proposed feature synthesis framework has been analyzed theoretically and empirically, and the resulting view-agnostic distance has been evaluated on various computer vision tasks, including fine-grained image retrieval and classification.

Future Work Our current framework does not take object occlusion and background clutter into consideration. We leave the task as a future work. In addition, the current surrogate region discovery method is at the categorical level, ignoring details of individual shapes. Geometric properties such as symmetry and part decomposition may help. Then one can use visible parts of an object to predict invisible parts.

References

[1] J. Acharya, A. Orlitsky, A. T. Suresh, and H. Tyagi. The complexity of estimating rényi entropy. *CoRR*, abs/1408.1000, 2014. 18, 19

[2] M. Aubry, D. Maturana, A. Efros, B. Russell, and J. Sivic. Seeing 3d chairs: exemplar part-based 2d-3d alignment using a large dataset of cad models. In *CVPR*, 2014. 11

[3] H. Bay, T. Tuytelaars, and L. Van Gool. Surf: Speeded up robust features. In *ECCV 2006*, pages 404–417. Springer, 2006. 2

[4] D.-Y. Chen, X.-P. Tian, Y.-T. Shen, and M. Ouhyoung. On visual similarity based 3d model retrieval. In *Computer graphics forum*, volume 22, pages 223–232. Wiley Online Library, 2003. 3

[5] J. Deng, A. C. Berg, and L. Fei-Fei. Hierarchical semantic indexing for large scale image retrieval. In *CVPR 2011*, pages 785–792. IEEE, 2011. 2

[6] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR 2009*, 2009. 10

[7] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR 2014*, 2014. 15

[8] Q.-X. Huang, H. Su, and L. Guibas. Fine-grained semi-supervised labeling of large shape collections. *ACM ToG*, 32(6):190:1–190:10, Nov. 2013. 10

[9] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014. 13

[10] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012. 2

[11] C. H. Lampert, H. Nickisch, and S. Harmeling. Learning to detect unseen object classes by between-class attribute transfer. In *CVPR 2009*, pages 951–958. IEEE, 2009. 2

[12] Q. V. Le, W. Y. Zou, S. Y. Yeung, and A. Y. Ng. Learning hierarchical invariant spatio-temporal features for action recognition with independent subspace analysis. In *CVPR 2011*, pages 3361–3368. 2

[13] D. Lowe. Object recognition from local scale-invariant features. In *ICCV*, 1999. 2

[14] S. Maji, E. Rahtu, J. Kannala, M. B. Blaschko, and A. Vedaldi. Fine-grained visual classification of aircraft. *CoRR*, abs/1306.5151, 2013. 15

[15] J. Matas, O. Chum, M. Urban, and T. Pajdla. Robust wide-baseline stereo from maximally stable extremal regions. *Image and vision computing*, 22(10):761–767, 2004. 2

[16] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000. 3

[17] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. S. Bernstein, A. C. Berg, and L. Fei-Fei. Imagenet large scale visual recognition challenge. *CoRR*, abs/1409.0575, 2014. 4

[18] S. Savarese and L. Fei-Fei. 3d generic object categorization, localization and pose estimation. In *ICCV*, 2007. 3

[19] H. Su, Q. Huang, N. J. Mitra, Y. Li, and L. Guibas. Estimating image depth using shape collections. *SIGGRAPH 2014*. 3

[20] J. B. Tenenbaum, V. De Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 2000. 3

[21] R. Vidal. A tutorial on subspace clustering. *IEEE Signal Processing Magazine*, 28(2):52–68, 2010. 8

6 Appendix

Please also refer to the .mp4 video at `http://ai.stanford.edu/~haosu/FeatureSynth/video_for_arxiv.mp4` for more visualization and explanation.

6.1 Intuition to Weight Transferrability (§3.4)

Assumption 1. *If a patch g_0 can perfectly surrogate g_1, and $\mathbf{x}_{v_0,g_0}$ and $\mathbf{x}_{v_1,g_1}$ are from some shape S such that $\mathbf{x}_{v_0,g_0} = \mathbf{S}_{:,v_0,g_0}\mathbf{w}_{v_0,g_0}$ and $\mathbf{x}_{v_1,g_1} = \mathbf{S}_{:,v_1,g_1}\mathbf{w}_{v_1,g_1}$, then $\mathbf{w}_{v_0,g_0} \equiv \mathbf{w}_{v_1,g_1}$.*

To provide more intuition, we check the simple cases when $\mathbf{x}_{v_0,g_0} = \mathbf{S}_{n,v_0,g_0}$ for $1 \leq n \leq N$, i.e., patch g_0 at view v_0 of shape $\mathbf{S}$ looks exactly like the corresponding patch of some $\mathbf{S}_n$. $\mathbf{x}_{v_0,g_0} = \mathbf{S}_{n,v_0,g_0}$ implies that $\mathbf{w}_{v_0,g_0} = [0, \cdots, 1, \cdots, 0]^T$, i.e., a zero vector except that the n-th element equals to 1[1]. On the other hand, because g_0 is a perfect patch surrogate of g_1, following the definition, $\mathbf{x}_{v_1,g_1} = \mathbf{S}_{n,v_1,g_1}$. This implies $\mathbf{w}_{v_1,g_1} = [0, \cdots, 1, \cdots, 0]^T$. Therefore, $\mathbf{w}_{v_1,g_1} = \mathbf{w}_{v_0,g_0}$.

[1]More rigorously, the result is true when the null space of $S_{:,v_0,g_0}$ is $\{\mathbf{0}\}$. In practice, the condition is usually satisfied when discriminative features such as HoG and DeepLearning features are used.

6.2 Derivation of Patch Surrogate Suitability (§3.2.1)

We first repeat the definition of surrogate suitability:

$$\begin{aligned}\gamma(g_0; g_1) &= \log P(A^i_{g_1} = A^j_{g_1} | A^i_{g_0} = A^j_{g_0}) \\ &= \log P(A^i_{g_0} = A^j_{g_0}, A^i_{g_1} = A^j_{g_1}) - \log P(A^i_{g_0} = A^j_{g_0})\end{aligned}$$

Please refer to §3.2.1 of main paper for notation definitions.

Next, we introduce the algorithm and sample complexity analysis for estimating $\gamma(g_0; g_1)$.

Lemma 1. $\gamma(g_0; g_1) = \log \sum_{(A_{g_0}, A_{g_1}) \in \mathcal{D} \times \mathcal{D}} P^2(A_{g_0}, A_{g_1}) - \log \sum_{A_{g_0} \in \mathcal{D}} P^2(A_{g_0})$

Proof. Computation of the first term:

$$\begin{aligned}&\log P(A^i_{g_0} = A^j_{g_0}, A^i_{g_1} = A^j_{g_1}) \\ =& \log \sum_{\substack{x \in \mathcal{D} \\ y \in \mathcal{D}}} P(A^i_{g_0} = x, A^j_{g_0} = x, A^i_{g_1} = y, A^j_{g_1} = y) \\ =& \log \sum_{(x,y) \in \mathcal{D} \times \mathcal{D}} P((A^i_{g_0}, A^i_{g_1}) = (x, y), (A^j_{g_0}, A^j_{g_1}) = (x, y)) \\ =& \log \sum_{(x,y) \in \mathcal{D} \times \mathcal{D}} P((A^i_{g_0}, A^i_{g_1}) = (x, y)) P((A^j_{g_0}, A^j_{g_1}) = (x, y)) \\ =& \log \sum_{(x,y) \in \mathcal{D} \times \mathcal{D}} P^2((A_{g_0}, A_{g_1}) = (x, y))\end{aligned}$$

where the last line follows from the independence of Shape i and Shape j.

Computation of the second term :

$$\begin{aligned}\log P(A^i_{g_0} = A^j_{g_0}) &= \log \sum_{x \in \mathcal{D}} P(A^i_{g_0} = x, A^j_{g_0} = x) \\ &= \log \sum_{x \in \mathcal{D}} P(A^i_{g_0} = x) P(A^j_{g_0} = x) \\ &= \log \sum_{x \in \mathcal{D}} P^2(A_{g_0} = x)\end{aligned}$$

where the last line follows from the independence of Shape i and Shape j. □

Note that, $\log \sum_x P^2(x)$ in Lemma 1 is a classical quantity in Information Theory, named Rényi-entropy. Recent work [1] provides a tight bound for the sample complexity of estimating Rényi-entropy. In the following, we restate their results.

Let f be an *estimator* of Rényi-entropy $\log \sum_x P^2(X)$ for distributions over support set $\mathcal{X}$, then $f : \mathcal{X}^* \to \mathbb{R}$ maps a sequence of samples drawn from a distribution p to its entropy. Given independent samples $X^n = X_1, \ldots, X_n$ from p, define

$$S^f(k, \delta, \epsilon) \stackrel{\text{def}}{=} \sup_p \{\min\{n : P(|H(p) - f(X^n)| > \delta) < \epsilon\}\}$$

where $k = |\mathcal{X}|$ is the cardinality of $\mathcal{X}$.

Then, the *sample complexity* of estimating $H(p)$ is defined as

$$S(k, \delta, \epsilon) \stackrel{\text{def}}{=} \inf_f S^f(k, \delta, \epsilon)$$

[1] shows that $S(k, \delta, \epsilon) = \Theta(\sqrt{k})$. The bound is achievable using the unbiased estimator.

Using the above result, we can prove the following result:

Theorem 2. *The optimal sample complexity to estimate $\gamma(g_0; g_1)$ is $\Theta(D)$, where D is the cardinality of symbol set $\mathcal{D}$ as defined in §3.2.1, i.e., the number of visual words in dictionary.*

Proof. It is easy to see that the sample complexity to estimate the first term of γ in Lemma 1 is $\Theta(D)$ and the second term is $\Theta(\sqrt{D})$. Therefore, the overall sample complexity for estimating γ is $\Theta(D)$. □

6.3 More Results on Image Retrieval

Next two pages show groups of representative image retrieval results using our view-agnostic distance and the baseline L2 HoG feature distance. For each group, the first image is the query image.

Figure 11: Image retrieval results on 100 classes (to be continued...)

Figure 12: Image retrieval results on 100 classes